

BOTTLENECK PARTIAL-MATCHING VORONOI DIAGRAMS AND APPLICATIONS

MATTHIAS HENZE AND RAFEL JAUME

ABSTRACT. Given two point sets in the plane, we study the minimization of the bottleneck distance between a point set B and an equally-sized subset of a point set A under translations. We relate this problem to a Voronoi-type diagram and derive polynomial bounds for its complexity that are optimal in the size of A . We devise efficient algorithms for the construction of such a diagram and its lexicographic variant, which generalize to higher dimensions. We use the diagram to find an optimal bottleneck matching under translations, to compute a connecting path of minimum bottleneck cost between two positions of B , and to determine the maximum bottleneck cost in a convex polygon.

1. INTRODUCTION

Applications often demand algorithms to find an occurrence of a point pattern in a given cloud of points. Using a suitable cost function, it is common to define a similarity measure between the pattern and the point cloud as the minimum cost among the images of the pattern under a set of allowed transformations. One of the most studied similarity measures between finite point sets A and B in \mathbb{R}^d is the *directed Hausdorff distance*, which is the maximum of the (Euclidean) distances from each point in B to its nearest neighbor in A . For some applications in robotics and pattern recognition, it is required that each point of the smaller set is matched to a distinct point in the bigger one. The resulting distance is called the *bottleneck distance* and was introduced for equally-sized sets in [2] as

$$\Delta(B, A) = \min_{\sigma: B \rightarrow A} \max_{b \in B} \|b - \sigma(b)\|,$$

where $\|\cdot\|$ denotes the Euclidean norm and the minimum is taken over all injections from B into A . In contrast to the directed Hausdorff distance, the bottleneck distance has the advantage of being symmetric for equally-sized sets. On the other hand, it is harder to compute, since the points cannot be regarded independently. Note that there might be several matchings that minimize the bottleneck distance, even when all the distances between points are distinct. When this is to be avoided, considering the matching that lexicographically minimizes the distances between matched points helps to break some ties; cf. [7, 11, 21].

Research by the first author is supported by the ESF EUROCORES program EuroGIGA-VORONOI, (DFG): RO 2338/5-1, and by the second author by “Obra Social la Caixa” and DAAD.

In this paper, we are interested in a dynamic version of the bottleneck distance. More precisely, we want to efficiently compute, among all translated copies of B with respect to A , one attaining the minimum bottleneck distance; that is, $\min_{t \in \mathbb{R}^d} \Delta(B + t, A)$. This problem will be called *bottleneck partial-matching under translations*. It was introduced for equally-sized point sets in the plane by Alt, Mehlhorn, Wagener & Welzl [2], who gave an algorithm running in $O(n^6 \log n)$ time for point sets of size n . Their bound was later improved to $O(n^5 \log^2 n)$ by Efrat, Itai & Katz [14].

To the best of our knowledge, bottleneck matching under translations has not been studied with the focus on algorithms whose complexity is sensitive to the size of the smaller set. In order to do so, we associate Voronoi-type diagrams to the problem, which we call *bottleneck diagrams* and *lex-bottleneck diagrams*, respectively. This follows an idea of Rote [20] who partitioned the space of translations according to the (partial) matching that minimizes the least-squares distance between translated copies of B and A (cf. [17, 4, 3] for follow-up studies). Our bottleneck diagrams partition \mathbb{R}^d into polyhedral cells that correspond to locally-optimal (lexicographic-)bottleneck matchings.

Our motivation to investigate these diagrams does not restrict to solving the bottleneck partial-matching problem under translations only. We moreover aim to provide a structure that may be either used for preprocessing or may be adjusted towards other problems that are based on the computation of the bottleneck distance in various translated positions of the point sets. The applications at the end of the paper exemplify this utility of the bottleneck diagrams.

A non-archival abstract containing parts of our studies appeared in [16].

Our Results. In Section 3, we formally introduce the Voronoi-type diagrams before investigating their basic properties and combinatorial complexity. It turns out that there exists a lex-bottleneck diagram (and, hence, a bottleneck diagram) of complexity $O(n^2 k^6)$ for any given planar point sets $A, B \subset \mathbb{R}^2$ with $k = |B| \leq |A| = n$ (see Theorem 3.12), and that this bound cannot be improved with respect to the size of A . For point sets $A, B \subset \mathbb{R}^d$ of higher dimensions we obtain in Corollary 3.6 that there is a lex-bottleneck diagram of complexity $O(n^{2d} k^{2d})$. Based on this complexity result, we devise algorithms in Section 4 that construct these polyhedral subdivisions of \mathbb{R}^d and at the same time compute a (lexicographic-)bottleneck matching for each of the cells of the subdivision, which is necessary to solve the bottleneck matching problem under translations. This is achieved with an overhead of $O(k^2)$ for the bottleneck diagrams, and $O(k^4)$ for the lexicographic variant (see Theorems 4.9 and 4.10). The matching problem under translations can then be solved for the bottleneck case in time $O(n^2 k^8)$, and for the lexicographic variant in time $O(n^2 k^{10})$, if the point sets are planar (see Theorem 5.1). In higher dimensions the time bounds are $O(n^{2d} k^{2d+2})$ and $O(n^{2d} k^{2d+4})$, respectively. Finally, Theorems 5.5 and 5.7 show how we can use the bottleneck diagrams to efficiently compute a path of minimum bottleneck cost between two positions of a pattern in the plane, or how to determine what we call the cover radius of a polygon.

Comparison to previous work. Although neither Alt et al. [2] nor Efrat et al. [14] consider the bottleneck matching problem for different-sized point sets, their methods can be adapted to this situation without major difficulties. In A, we elaborate on such an analysis of their algorithms and derive the time bounds $O(n^3 k^3 \log n)$ and $O(n^2 k^3 \log^2 n)$, respectively, where k is the size of the smaller set and n the size of the bigger one. This shows that the adapted algorithm by Efrat et al. outperforms our procedure, that runs in $O(n^2 k^8)$, already for fairly small values of k , in particular for any $k = \Omega(\log^{2/5} n)$. Still, the use of bottleneck diagrams is conceptually different from previous methods, and has the advantage of being applicable to solve the translative matching problem in any dimension and, moreover, with respect to the lexicographic bottleneck cost. No exact algorithms were known for higher dimensions previously, however, there exist approximation algorithms for the bottleneck matching problem (see [14]).

Organization of the paper. In the next section, we introduce (lexicographic) bottleneck matchings between two finite point sets and investigate corresponding minimization diagrams and their properties. After these basics, we define our main objects of study, bottleneck partial-matching Voronoi diagrams, and analyze their combinatorial complexity in Section 3, before addressing construction algorithms for these structures in Section 4. Finally, in Section 5, we apply the bottleneck diagrams to solve the bottleneck partial-matching problem and related questions.

2. BOTTLENECK AND LEXICOGRAPHIC BOTTLENECK MATCHINGS

In this section, we introduce bottleneck matchings and discuss the minimization diagram corresponding to the bottleneck partial-matching problem under translation. The issues we encounter explain our approach to the definition of the bottleneck diagrams in Section 3.

Throughout the paper, we assume that we are given two point sets $A, B \subset \mathbb{R}^d$ with $k = |B| \leq |A| = n$ and that B is allowed to be translated. We use the term *edge* for a pair of points $(a, b) \in A \times B$ and denote it by ab for short. The *length* of the edge ab is defined as the Euclidean distance $\|b - a\|$. In this context, we identify every injection of B into A with the *matching*, i.e., the set of edges, it induces. The cost of such a matching varies according to a parameter representing the position of the point set B .

Definition 2.1. Let $A, B \subset \mathbb{R}^d$ be finite point sets with $|B| \leq |A|$. A *bottleneck matching* for A and B is a matching that minimizes

$$f(\sigma) = \max_{b \in B} \|b - \sigma(b)\| \text{ among all the matchings } \sigma : B \hookrightarrow A.$$

The *bottleneck cost* of a matching $\sigma : B \hookrightarrow A$ is the function $f_\sigma : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ defined as

$$f_\sigma(t) = \max_{b \in B} \|b + t - \sigma(b)\|^2, \text{ for all } t \in \mathbb{R}^d.$$

The *bottleneck value function* $\mathcal{E} : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is defined by

$$\mathcal{E}(t) = \min_{\sigma : B \hookrightarrow A} f_\sigma(t), \text{ for all } t \in \mathbb{R}^d.$$

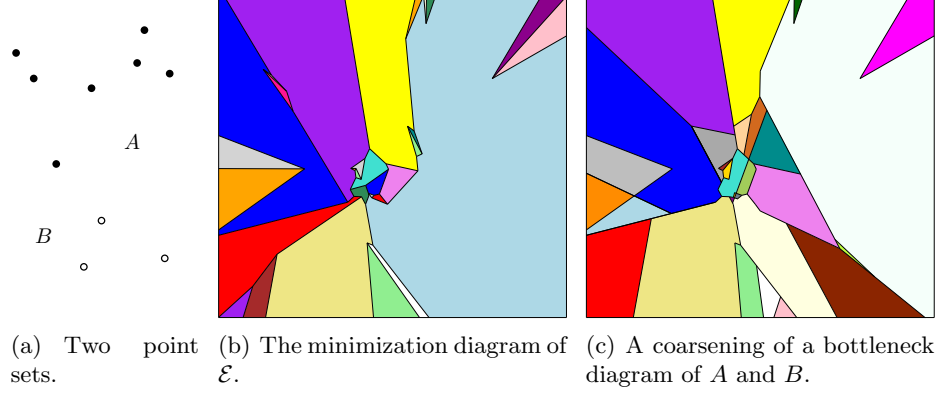


FIGURE 1. The minimization diagram of \mathcal{E} and a coarsening of a bottleneck diagram for a pair of point sets.




FIGURE 2. Three positions of a pair of point sets and a bottleneck matching for each of them showing that the minimization diagram of the corresponding \mathcal{E} function has a non-convex region.

Note that a bottleneck matching is defined in terms of the Euclidean distance while the functions f_σ and \mathcal{E} depend on the square of this value. This squaring is harmless and will be convenient later on.

By definition, the function \mathcal{E} is piecewise quadratic and induces a partition of \mathbb{R}^d into polyhedral regions which is usually called the *minimization diagram* of \mathcal{E} . Figs. 1(a) and 1(b) show a pair of planar point sets and the corresponding minimization diagram whose regions are colored according to the edges attaining the bottleneck value. Note that some regions are not convex and some are even disconnected. More precisely, the red and the blue regions consist of two connected components. The possible non-convexity of these regions is more concisely illustrated in Fig. 2. The pictured disks certify that the drawn edges are the ones attaining the bottleneck value for the three aligned positions of the small point set.

Moreover, the leftmost and the rightmost matchings in Fig. 2 are the only optimal matchings for the respective positions of the small point set. On the other hand, they have the same longest edge, i.e., the quadratic functions in the corresponding regions of the minimization diagram coincide. Note in addition that, disregarding the longest edge, these two matchings have disjoint sets of edges. Conversely, it can happen that a matching is the unique bottleneck matching for two open sets of positions contained in different regions of the minimization diagram of \mathcal{E} . An instance of this

situation is illustrated in Fig. 3; remember that different longest edges in the matching correspond to different regions.



FIGURE 3. A bottleneck matching with different longest edges in different positions of B .

Another observation is that, even for a fixed position of the point sets, there are in general many bottleneck matchings. Fig. 4 shows an example that can be easily generalized to show that a set of k points and a set of $n \geq k$ points can have $(n-1)/(n-k)!$ different bottleneck matchings. Indeed, in the depicted situation only the edge between the points b_0 and a_0 is fixed for a matching to be bottleneck, the remaining points can be matched arbitrarily.




FIGURE 4. A pair of point sets with many bottleneck matchings.

Moreover, the bottleneck matchings for the point sets in Fig. 4 all have the same longest edge, and they remain bottleneck matchings if one of the point sets is translated in an open neighborhood of its current position. However, there may be different edges that are the longest edge of bottleneck matchings everywhere in a neighborhood of a fixed position, as illustrated in Fig. 5. Just like in the previous figures, the bold edges attain the bottleneck value.

A common way to break ties between bottleneck matchings for a given position in order to be more sensitive to the geometry of the point sets is to consider a lexicographic version of bottleneck matchings. That is, among the matchings whose longest edge is as short as possible, consider those whose second longest edge is as short as possible, and so on. For the precise definition, we recall that the *lexicographic order* on \mathbb{R}^k is the total order induced by the relation $(x_1, \dots, x_k) \prec (y_1, \dots, y_k)$ if and only if there exists an $m \in [k]$ such that $x_i = y_i$ for all $i < m$, and $x_m < y_m$. We write $x \preceq y$ if $x \prec y$ or $x = y$.

Definition 2.2. Let $A, B \subset \mathbb{R}^d$ be two finite point sets with $k = |B| \leq |A|$. The *lex-bottleneck cost* of a matching $\sigma : B \hookrightarrow A$ is the function $g_\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^k$ where the i -th coordinate of $g_\sigma(t)$ corresponds to the length of the i -th longest edge of σ for $B + t$ and A . A *lex-bottleneck matching* for A and B is a matching π such that $g_\pi(0) \preceq g_\sigma(0)$ for every other matching σ .




FIGURE 5. Bottleneck matchings with different longest edges.

Note that a lex-bottleneck matching is a bottleneck matching as well. Although this definition certainly breaks some ties, we see in the next section that it does not guarantee uniqueness (not even in the complement of a nowhere-dense set).

3. BOTTLENECK PARTIAL-MATCHING VORONOI DIAGRAMS

In this part, we define Voronoi-type diagrams associated with the (lexicographic) bottleneck partial-matching problem under translation. We discuss basic properties of these structures and derive reasonable bounds on their combinatorial complexity.

3.1. Definitions and Basic Properties. As we have seen in the previous section, we face some difficulties when we define a Voronoi-type structure for the bottleneck partial-matching problem. One of them is the existence of open sets of translations for which neither the bottleneck matching nor the longest edge are uniquely determined. The problem of the non-uniqueness of the longest edge can be solved by requiring the point set to be in an ad hoc general position. The non-uniqueness of the matching may be attacked by considering the lexicographic variant, but even in general position it might fail even in open balls. Nevertheless, it is of interest to study the original version as well in order to solve problems like the ones in Section 5 or to explore the minimization diagram of \mathcal{E} .

As expected, the Voronoi-type diagrams we are going to study can be required to be given in form of a polyhedral complex. This facilitates traversing the partition or optimizing in a region, operations that are often required in related problems. In what follows, full-dimensional faces of a polyhedral complex are called *cells*.

Definition 3.1. Let $A, B \subset \mathbb{R}^d$ be finite point sets with $|B| \leq |A|$. A *bottleneck partial-matching Voronoi diagram* (or *bottleneck diagram*, for short) for A and B is a polyhedral complex \mathcal{T} covering \mathbb{R}^d and such that for every cell C of \mathcal{T} there is at least one matching $\pi_C : B \hookrightarrow A$ such that $f_{\pi_C}(t) \leq f_{\sigma}(t)$ for all $t \in C$ and all matchings $\sigma : B \hookrightarrow A$. A *bottleneck labeling* of a bottleneck diagram is a function mapping each cell to one such matching.

A coarsening of a bottleneck diagram of the point set in Fig. 1(a) is displayed in Fig. 1(c), where cells with the same label have the same color. Note that for $B = \{b\}$ the Voronoi diagram of $A - b$ is a bottleneck diagram. A diagram for the lexicographic version of the bottleneck cost is defined analogously.

Definition 3.2. Let $A, B \subset \mathbb{R}^d$ be two finite point sets with $|B| \leq |A|$. A *lex-bottleneck partial-matching Voronoi diagram* (or *lex-bottleneck diagram*, for short) for A and B is a polyhedral complex \mathcal{T} covering \mathbb{R}^d and such that for every face c of \mathcal{T} there is at least one matching $\pi_c : B \hookrightarrow A$ such that $g_{\pi_c}(t) \preceq g_{\sigma}(t)$ for all t interior to c and all matchings $\sigma : B \hookrightarrow A$. A *lex-bottleneck labeling* of a lex-bottleneck diagram is a function mapping each face to one such matching.




FIGURE 6. A pair of point sets and a lex-bottleneck labeling of the neighborhood of the represented position.

Note that a bottleneck matching in a lower-dimensional face of a bottleneck diagram is given by the labeling of a cell containing it, i.e., the label for a cell is valid everywhere in the cell. This is not the case for a lex-bottleneck diagram, which is the reason why a label for each face is required. Fig. 6 shows an example. On the left of the figure, the point sets A and B are displayed, for which the blue and the red matchings are both lex-bottleneck matchings. On the right, a small neighborhood in a lex-bottleneck diagram around the point $t \in \mathbb{R}^2$ corresponding to the depicted position of the point sets is represented. If the point set B (the white dots) is moved infinitesimally to the right, only the blue matching is a lex-bottleneck matching, whereas if it is moved infinitesimally to the left, only the red matching remains lex-bottleneck. This forces the cyan and orange regions to be labeled with the blue and red matchings, respectively. However, if B is vertically translated by an infinitesimal amount, the longest blue edge and the longest red edge have the same length. In addition, the second longest red edge is longer than the second longest blue edge if the perturbation is upwards, while it is shorter if the perturbation is downwards, forcing the blue and red regions to be labeled accordingly. For the depicted position (corresponding to the purple point), both matchings are equally good, since the respective shortest edges are equally long.

Since any lex-bottleneck diagram is a bottleneck diagram, we prove some properties for the first, more restrictive type. However, later on we devise algorithms that compute a bottleneck labeling more efficiently than a lex-bottleneck one. For our applications in Section 5 the first type of labeling is enough.




FIGURE 7. Two lex-bottleneck matchings, represented for three positions of B .

Definition 3.3. Given $x, y, v, z \in \mathbb{R}^d$, let

$$\begin{aligned} h(x, y, v, z) &= \left\{ t \in \mathbb{R}^d : \|y + t - x\| = \|z + t - v\| \right\} \\ &= \left\{ t \in \mathbb{R}^d : 2 \langle t, y - x - (z - v) \rangle = \|z - v\|^2 - \|y - x\|^2 \right\}. \end{aligned}$$

Given two finite point set $A, B \subset \mathbb{R}^d$, let $\mathcal{H}(A, B)$ be the arrangement of the hyperplanes $h(a, b, a', b')$, called *bisectors*, for all pairs $a, a' \in A$ and $b, b' \in B$ such that $b - a \neq b' - a'$.

Proposition 3.4. *The arrangement $\mathcal{H}(A, B)$ is a lex-bottleneck diagram.*

Proof. The squared length of an edge matching b to $a = \sigma(b)$ is given by

$$\|b + t - a\|^2 = \|t\|^2 + \|b - a\|^2 + 2 \langle t, b - a \rangle.$$

For a pair of edges $ab, a'b' \in A \times B$, the locus of points $t \in \mathbb{R}^d$ for which $\|b + t - a\|^2 = \|b' + t - a'\|^2$ is exactly $h(a, b, a', b')$. If $b - a \neq b' - a'$, then this set is a hyperplane, otherwise the whole \mathbb{R}^d . Let c be a face in $\mathcal{H}(A, B)$, and let π be a lex-bottleneck matching for a point t_0 in the relative interior of c . A lex-bottleneck matching for the translation $t \in \mathbb{R}^d$ only depends on the relative length of the possible edges in $A \times (B + t)$. Hence, the matching π remains lex-bottleneck for t as well as long as no edge becomes strictly shorter than another edge that was strictly longer for t_0 . By definition of $\mathcal{H}(A, B)$, this cannot happen in the relative interior of c . Since an arrangement of hyperplanes is a polyhedral complex, $\mathcal{H}(A, B)$ is a lex-bottleneck diagram for A and B . \square

Observe that there may be open sets for which two different matchings are lex-bottleneck matchings, as shown in Fig. 7. This is because every edge from the red matching can be paired with an edge of the blue matching having the same length for any position of the matching. As long as the blue match and the red match of every point in B are its two closest points (as in the three positions represented in the figure), both are lex-bottleneck matchings. Nonetheless, we see now that the *matched sets* $\sigma(B)$ of such matchings $\sigma : B \hookrightarrow A$ coincide.

Proposition 3.5. *Let $A, B \subset \mathbb{R}^d$ be finite point sets with $|B| \leq |A|$. If two matchings are lex-bottleneck matchings in an open set $U \subset \mathbb{R}^d$, then they have the same matched set and they have the same lex-bottleneck cost for any $t \in \mathbb{R}^d$. In particular, in the interior of a cell of $\mathcal{H}(A, B)$ there is a*

unique subset $A' \subseteq A$ that is the matched set of all lex-bottleneck matchings for that cell.

Proof. As argued in the proof of Proposition 3.4, the set of lex-bottleneck matchings is the same for all the translations interior to a cell of $\mathcal{H}(A, B)$. Let C be a cell of $\mathcal{H}(A, B)$, t_0 be a point interior to C , and π and σ be two different lex-bottleneck matchings for $B + t_0$ and A . We regard here the matchings π and σ as sets of edges on the vertex set $A \cup B$. Consider then the symmetric difference of π and σ , i.e., the set of edges that belong to π or to σ but not to both. The graph with this set of edges is a collection of (even length) vertex-disjoint paths and cycles, whose edges from π alternate with edges from σ , since the edges of a matching are pairwise-disjoint (see Section 4.1 for details).

In addition, since we assume that both matchings are lex-bottleneck matchings, we have that there is a one-to-one correspondence between edges of every path or cycle belonging to π and the edges from the same path or cycle belonging to σ , such that the corresponding edges have the same length. Indeed, assume for a contradiction that there is no such correspondence for some path or cycle γ . Then, the restriction of one of the matchings, say π , to γ would be better than the restriction of the other (in the lexicographic sense). The result of replacing in σ the edges of $\sigma \cap \gamma$ with the edges of $\pi \cap \gamma$ is a matching and it is lexicographically better than σ , which contradicts its assumed optimality.

It follows from the proof of Proposition 3.4, that the edges ab and $a'b'$ have the same length over an open set if and only if $b - a = b' - a'$. Thus, there is no path in the symmetric difference, because following the edges in a path we would arrive to the starting vertex, since every edge is “cancelled” by the corresponding edge of the other matching. Therefore, the symmetric difference is made exclusively of cycles, and hence the matched sets of π and σ coincide. \square

3.2. Complexity of the Bottleneck Diagrams. The construction in the proof of Proposition 3.4 leads immediately to a first bound on the complexity of a lex-bottleneck diagram.

Corollary 3.6. *For any pair of point sets $A, B \subset \mathbb{R}^d$ with $k = |B| \leq |A| = n$, there is a lex-bottleneck diagram of combinatorial complexity $O(n^{2d}k^{2d})$.*

Proof. By Proposition 3.4 the arrangement $\mathcal{H}(A, B)$ is a lex-bottleneck diagram for A and B . It is well-known (cf. [12]) that the complexity of an arrangement of m hyperplanes in \mathbb{R}^d is $O(m^d)$. The arrangement $\mathcal{H}(A, B)$ consists of $\binom{n}{2}\binom{k}{2} = O(n^2k^2)$ hyperplanes in \mathbb{R}^d , and hence the claimed bound follows. \square

Our aim is now to show that there are many hyperplanes of $\mathcal{H}(A, B)$ that we can safely ignore. This allows for an improvement of the above bound for planar point sets. We need to introduce some necessary notation and collect some auxiliary results before we can present our argument.

First of all, we state some simple results on hyperplane arrangements, the first of which can be found, e.g., in [12], and will be used implicitly later on.

Proposition 3.7. *For a finite point set $P \subset \mathbb{R}^d$, consider the linear functions*

$$h_p(x) = \|p\|^2 - 2\langle x, p \rangle, \text{ for } p \in P.$$

For a fixed $k \in [|P|]$, the k -level of the arrangement of hyperplanes associated to these functions projects onto the order- k Voronoi diagram of P .

Proposition 3.8. *Let $S \subset \mathbb{R}^2$ be a finite point set and let \mathcal{Z} be the set of planes*

$$z_p = \{(x, y, \|p\|^2 - 2\langle (x, y), p \rangle) : (x, y) \in \mathbb{R}^2\} \subset \mathbb{R}^3, \text{ for } p \in S.$$

Every three planes in \mathcal{Z} intersect in at most one point. Equivalently, the locus of points equidistant from three different points of S is either a point or empty.

Proof. For every point $q \in \mathbb{R}^2$, the vertical order of the planes z_p over q is the same as the distances from q to the corresponding $p \in S$. The locus of points equidistant from three points in the plane is the center of the circle through them, or the empty set if they are collinear. Hence, three planes can coincide in at most one point. \square

The main ingredient to improve the complexity bound in Corollary 3.6 is the following lemma.

Lemma 3.9 (Ben-Avraham et al. [4]). *Let $A \subset \mathbb{R}^2$ be a set of n points. There are $O(nk)$ bisectors that support all edges of order- j Voronoi diagrams of A for all $j \leq k$.*

Finally, we fix some technicalities for lex-bottleneck matchings.

Definition 3.10. Let $A, B \subset \mathbb{R}^d$ be two finite point sets with $k = |B| \leq |A|$. Given $b \in B$, a set of edges $M_b \subseteq A \times B$ is called a b -minimal set if $|M_b| = k$ and no edge $ab \in (A \times B) \setminus M_b$ is strictly shorter than any edge in M_b . The candidate set of $b \in B$ at a position $t \in \mathbb{R}^d$ is the set $E_b(t) \subseteq A \times (B + t)$ with $|E_b(t)| \geq k$ and such that every edge $(a, b + t) \notin E_b(t)$ is strictly longer than every edge in $E_b(t)$. Furthermore, we write $Z(t) = \cup_{b \in B} E_b(t)$.

Lemma 3.11. *Let $A, B \subset \mathbb{R}^d$ be two finite point sets with $k = |B| \leq |A|$.*

- (i) *If a subset $M \subseteq A \times B$ contains a b -minimal set for each $b \in B$, then M contains a lex-bottleneck matching for A and B .*
- (ii) *Every lex-bottleneck matching for A and B is contained in the union $Z = Z(0)$ of candidate sets.*

Proof. (i): Let $\mu \subset A \times B$ be a lex-bottleneck matching. If $\mu \subseteq M$, nothing is left to prove. Otherwise, let $ab \in A \times B$ be an edge in $\mu \setminus M$, and let $M_b \subseteq M$ be a b -minimal set. Since μ matches exactly k points of A and $ab \notin M \supseteq M_b$, there must be an edge $a'b \in M_b \setminus \mu$. The matching $(\mu \setminus \{ab\}) \cup \{a'b\}$ is lexicographically at least as good as μ and it uses one more edge of M . Repeating the process, we end up with a lex-bottleneck matching contained in M .

(ii): Assume that μ is a lex-bottleneck matching for A and B and that it contains an edge $ab \notin Z$. Let $a'b$ be an edge in $E_b(0) \setminus \mu \neq \emptyset$. By definition of candidate sets, the matching $(\mu \setminus \{ab\}) \cup \{a'b\}$ is lexicographically better than μ , contradicting its optimality. \square

Now we are well-prepared to prove our main result of this section.

Theorem 3.12. *For any pair of point sets $A, B \subset \mathbb{R}^2$ with $k = |B| \leq |A| = n$, there is a lex-bottleneck diagram of complexity $O(n^2 k^6)$.*

Proof. Observe that $E_b(t_1) = E_b(t_2)$ for every pair of points t_1, t_2 interior to a cell C of $\mathcal{H}(A, B)$, for all $b \in B$. Therefore, we write $E_b(C) = E_b(t)$ and $Z(C) = Z(t)$ for any point t interior to a cell C of $\mathcal{H}(A, B)$. By Lemma 3.11-(ii), any lex-bottleneck matching for a point interior to C is contained in $Z(C)$. Let now t_0 be a point in a lower-dimensional face L of $\mathcal{H}(A, B)$, and let $C \supset L$ be a cell. For continuity reasons, the set $E_b(C) \subseteq E_b(t_0)$ is a b -minimal set at t_0 as well. Therefore, by Lemma 3.11-(i), there is a lex-bottleneck matching for t_0 contained in $Z(C)$.

Now, consider a labeling Λ of $\mathcal{H}(A, B)$ that labels every face of a cell C with a matching contained in $Z(C)$. We say that an edge F of $\mathcal{H}(A, B)$ between two cells C_l and C_r *uses* a bisector $h = h(a, b, a', b')$ if $F \subset h$ and $ab, a'b' \in Z(C_l) \cup Z(C_r)$. If no edge uses a bisector h , then h can be omitted from the arrangement $\mathcal{H}(A, B)$ while maintaining its property of being a lex-bottleneck diagram. Each new face in the resulting hyperplane arrangement is a union of a set of old faces, i.e., before removing h , and can be labeled according to the label of any of them, since h does not intersect any face whose label in Λ uses the edges ab or $a'b'$. We show now that many bisectors are not used by any edge, distinguishing the two following cases.

Let $h(a, b, a', b)$ be a bisector used by an edge $F = C_l \cap C_r$, and consider the point set $S(b) = A - b$. The edge F must be contained in an order- j Voronoi edge of $S(b)$ for some $j \leq k$. Indeed, in view of Proposition 3.8, infinitesimally to the right or to the left of (a point in the relative interior of) F both $a - b$ and $a' - b$ are among the $k + 1$ closest points. In addition, for a point t_0 in the relative interior of F , the points $a - b$ and $a' - b$ are the only two points of $S(b)$ that lie at distance $\|b + t_0 - a\| = \|b + t_0 - a'\|$. Hence, there is a circle centered at t_0 and through $a - b$ and $a' - b$ that contains $j - 1 \leq k - 1$ points in its interior, which is a characterization for points in the relative interior of edges of the order- j Voronoi diagram.

Let $h(a, b, a', b')$ with $b \neq b'$ be a bisector used by an edge $F = C_l \cap C_r$, and consider the point set $S(b, b') = (A - b) \cup (A - b')$. Note that the number of points in $S(b, b')$ is not necessarily $2n$: a point $a_1 - b = a_2 - b'$ can belong to both $A - b$ and $A - b'$. However, this is the case if and only if the edges $a_1 b$ and $a_2 b'$ are equally long everywhere. In particular, the points $a - b$ and $a' - b'$ are distinct since, otherwise, they would not induce any bisector. Furthermore, simple algebraic manipulations show that $t \in \mathbb{R}^2$ is closer to $a_1 - b$ than to $a_2 - b'$ if and only if $b + t$ is closer to a_1 than $b' + t$ is to a_2 , for any choice of $a_1, a_2 \in A$. Since the bisector is used by F , the point $a - b$ is among the k closest points of $A - b$ infinitesimally to at least one of the sides of F . Proposition 3.8 ensures that along the interior of F only the points $a - b$ and $a' - b'$ are at distance $\|b + t_0 - a\| = \|b' + t_0 - a'\|$ among the points in $S(b, b')$, which implies that in fact $a - b$ is among the k closest points of $A - b$ infinitesimally to both sides of F . Similarly, the point $a' - b'$ belongs to the k closest points of $A - b'$ for points infinitesimally away from F . Hence, for any point t_0 in the relative interior of F , there is a disk centered at t_0 and passing through $a - b$ and $a' - b'$ that contains $j - 1 \leq 2k - 2$ points

of $S(b, b')$ in its interior. Equivalently, the bisector $h(a, b, a', b')$ supports an order- j Voronoi edge of $S(b, b')$ for some $j \leq 2k - 1$.

Applying Lemma 3.9 to $S(b)$, for all $b \in B$, and to $S(b, b')$, for every pair $b, b' \in B$, it follows that the number of bisectors that are used by some edge is $O(k^2 \cdot nk)$. The complexity of the diagram resulting from removing all unused bisectors from $\mathcal{H}(A, B)$ is thus $O(n^2 k^6)$. \square

Based on a one-dimensional example of Rote [20], one may derive a lower bound on the complexity of any lex-bottleneck diagram. The proof is a particular case of the result for stable matchings that can be found in [3].

Proposition 3.13. *For any $k, n \in \mathbb{N}$ with $n \geq k \geq d$, there exist point sets $A, B \subset \mathbb{R}^d$ with $|B| = k$ and $|A| = n$ such that any lex-bottleneck diagram for A and B has complexity $\Omega(k^d(n - k)^d)$.*

In particular, this shows that Theorem 3.12 is optimal with respect to the size of the bigger set A .

4. CONSTRUCTION OF THE BOTTLENECK DIAGRAMS

In this section we are concerned with construction algorithms for the bottleneck diagrams introduced before. The basic idea is to first construct the reduced hyperplane arrangement discussed in the proof of Theorem 3.12, and then traverse it while computing a bottleneck labeling for each cell (or face) of the structure. The latter is based on the well-developed theory of bottleneck assignments in weighted bipartite graphs whose concepts and methods that are relevant for our purposes we introduce first.

4.1. Notation and Techniques for Matchings in Weighted Bipartite Graphs. As usual, let $A, B \subset \mathbb{R}^d$ be finite point sets with $|B| \leq |A|$. The problem of finding a bottleneck matching for a fixed position of B can be translated into a matching problem in a weighted bipartite graph on A and B , where the weight of an edge from $ab \in A \times B$ is the Euclidean distance between the corresponding points. Most of the geometric definitions from Sections 2 and 3 have graph-theoretic analogs. Note that we use the same symbols and names for corresponding concepts in the graph setting.

Let $G = (A, B; E)$ be a bipartite graph with edge set $E \subseteq A \times B$ and vertex set partitioned into the components A and B . A *matching* in G is a set $\sigma \subseteq E$ such that every vertex in $A \cup B$ is incident to at most one edge of σ . As in the geometric setting, we identify a matching with the injection from B into A it induces, and we simplify notation by denoting an edge $(a, b) \in E$ by ab . A *maximum matching* is a matching of maximum cardinality. Vertices that belong to an edge of a matching are called *matched vertices* and otherwise *exposed vertices*.

An *alternating path/cycle* for a matching σ is a path/cycle in G with no repeated vertices such that the even edges are in σ and the odd ones are in $E \setminus \sigma$. An *augmenting path* for σ is an alternating path starting and ending at exposed vertices. Note that, if γ is an augmenting path for σ , the matching $\sigma \oplus \gamma = (\sigma \setminus \gamma) \cup (\gamma \setminus \sigma)$ has one more matched vertex than σ . In general, for two sets of edges $\sigma, \tau \subseteq E$ the connected components of the graph induced by $\sigma \oplus \tau$ are called its *components*. If σ and τ are maximum

matchings, the components of $\sigma \oplus \tau$ are paths or cycles, since every vertex has degree at most two in $\sigma \oplus \tau$.

Bottleneck Assignments in Bipartite Graphs. Given a bipartite graph $G = (A, B; E)$, we let $w : E \rightarrow \mathbb{R}_{\geq 0}$ be a function giving weights to its edges. The bottleneck cost of a matching π in G with respect to w is the maximum w -value attained by the edges in π . The problem of finding a maximum matching of minimum bottleneck cost, henceforth referred to as a *bottleneck matching*, for the complete and balanced case, i.e., $E = A \times B$ and $k = n$, has been widely studied in the last decades under the name of the *bottleneck assignment problem*. The most prominent approaches for this problem are the *threshold methods* and the *augmenting path methods*. Details and related studies can be found in the book dedicated to assignment problems from Burkard, Dell’Amico & Martello [6]. The threshold algorithms conduct a binary search on the possible values for the edge with maximum weight of a bottleneck matching. At each stage, the edges with bigger weight than the threshold are ignored, and a maximum matching computation in the modified graph is performed. One of the best-known algorithms to find a maximum matching in a bipartite graph is due to Hopcroft & Karp [18]. It runs in $O(|E|\sqrt{\nu(G)})$ time, where $\nu(G)$ is the size of the maximum matchings in G . The algorithm by Alt, Blum, Mehlhorn & Paul [1] finds a maximum matching more efficiently if the graph is “dense”. This fact was exploited by Punnen & Nair [19] to develop an alternative algorithm for the bottleneck assignment problem. The pure threshold method is preferable for dense graphs, the method from Gabow & Tarjan [15] is better for sparse graphs and the approach of Punnen & Nair covers the range in between.

Similar to the geometric situation, a variant of the bottleneck assignment problem is the *lexicographic bottleneck assignment problem*, introduced in [7] and revisited in [21]. For this problem, the cost of a matching $\sigma : B \hookrightarrow A$ is the result of sorting decreasingly the values $w(\sigma(b)b)$, for all $b \in B$. A *lexicographic bottleneck matching* is a matching minimizing the cost, when the corresponding cost vectors are compared lexicographically. Note that a lexicographic bottleneck matching is necessarily a bottleneck matching.

In the unbalanced case, as long as the graph is complete, the maximum matchings have size $|B|$. Let us recall the concepts of b -minimal and candidate sets from Definition 3.10 by reviewing their definition in the graph setting.

Definition 4.1. Let $G = (A, B; E)$ be a bipartite graph with $k = |B| \leq |A|$, and $w : E \rightarrow \mathbb{R}_{\geq 0}$ be a function giving weights to the edges. Given $b \in B$, a set $M_b \subseteq E$ is called a *b -minimal set* if $|M_b| = k$ and no edge $ab \in E \setminus M_b$ has strictly smaller weight than any edge in M_b . The *candidate set of b* is the set $E_b \subseteq E$ with $|E_b| \geq k$ and such that every edge $ab \in E \setminus E_b$ has strictly larger weight than any edge in E_b .

Lemma 3.11 holds analogously for bipartite graphs and can be considered as a property similar to Hall’s marriage theorem. As a consequence thereof, we can select a set of k^2 edges which is guaranteed to contain a lexicographic bottleneck matching. This pruning of the graph can be done in $O(nk)$ time using selection algorithms. Although we do not know whether the graph will

be dense or sparse after pruning the non-relevant edges and isolated vertices, we have that both $|B| + |A|$ and $|E|$ are $O(k^2)$. Thus, the best worst-case running time for our scenario is provided by the algorithm of Gabow & Tarjan, which runs in $O(k^2\sqrt{k}\log k)$ time, according to the analysis in [6]. The approach in [21], based on solving a sequence of linear sum assignment problems and bottleneck assignment problems, boils down to an algorithm for the computation of a lexicographic bottleneck matching running in $O(k^4)$ time after the pruning.

We close this part by stating some well-known results in the matching literature, which are used by the aforementioned algorithms and required below.

Lemma 4.2 (Berge’s Lemma [5]). *A matching σ is a maximum matching in a bipartite graph G if and only if there is no augmenting path for σ in G .*

Definition 4.3. Let $G = (A, B; E)$ be a bipartite graph and $w : E \rightarrow \mathbb{R}_{\geq 0}$ be a function giving weights to its edges. Given $r \in [|E|]$, let

$$E(r) = \{e \in E : w(e) \text{ is among the } r \text{ smallest values of } w(E)\}.$$

We denote by $G(r)$ the graph $G = (A, B; E(r))$.

We say that a matching in a bipartite graph is a *complete matching* if it matches all the points in the smaller set of vertices.

Proposition 4.4. *Let $G = (A, B; E)$ be a bipartite graph that has a complete matching and let $w : E \rightarrow \mathbb{R}_{\geq 0}$ be a function giving weights to the edges of G .*

- (i) *If $G(r)$ has a complete matching, then $G(j)$ has a complete matching for all $j > r$.*
- (ii) *A complete matching in $G(r)$ is a bottleneck matching if and only if $G(r-1)$ has no complete matching.*

4.2. Main Construction Theorem. In this section, we discuss some algorithmic techniques in order to construct a labeled bottleneck and a lex-bottleneck diagram for a pair of point sets in the plane. To this end, we introduce the following notation.

Definition 4.5. For finite point sets $A, B \subset \mathbb{R}^2$, we denote by $\mathcal{L}(A, B)$ the arrangement constructed in the course of the proof of Theorem 3.12. Given a cell C of $\mathcal{L}(A, B)$, we denote by $Z(C)$ the set $Z(t)$ for any point t interior to C .

We first describe how to use established techniques in order to construct an unlabeled lex-bottleneck diagram.

Lemma 4.6. *The lex-bottleneck diagram $\mathcal{L}(A, B)$ for point sets $A, B \subset \mathbb{R}^2$ with $k = |B| \leq |A| = n$ can be constructed in $O(n^2k^6)$ time.*

Proof. An arrangement of m lines in the plane can be computed in $O(m^2)$ time using an optimal algorithm, such as the incremental algorithm [9] or a topological sweep [13]. However, the proof of Theorem 3.12 is not constructive and, hence, it is not obvious how to select the bisectors that are used by some edge (in the sense of Theorem 3.12) among the $O(n^2k^2)$ candidates. Fortunately, the algorithm by Chan [8] constructs the facial structure

of the $(\leq s)$ -level of an arrangement of m planes in $O(m \log m + ms^2)$ expected time. In addition, this algorithm can be derandomized, leading to a deterministic version running in $O(ms^2(\log m / \log s)^{O(1)})$ time. We can then construct the $O(k^2)$ necessary structures in $O(k^4 n (\log n / \log k)^{O(1)})$ and traverse each of them discovering the $O(nk^3)$ used bisectors to finally construct their arrangement. \square

We now show how to find a bottleneck labeling of $\mathcal{L}(A, B)$. Before detailing the algorithm, we need a technical lemma that examines how small changes in a graph affect its bottleneck matchings. Recall that a matching in a bipartite graph on A and B is a complete matching if it matches all the points of B .

Lemma 4.7. *Let $G = (A, B; E, w)$ be a bipartite graph with $w : E \rightarrow [|E|]$ giving weights to its edges. Let μ be a bottleneck matching for G , and let $l \in E$ be the longest edge of μ in G . For a fixed $j \in [|E| - 1]$, let $G' = (A, B; E, w')$ be the weighted bipartite graph, where w' coincides with w except that $w'(e) = j$ if $w(e) = j + 1$, and $w'(e) = j + 1$ if $w(e) = j$, for all edges $e \in E$.*

- (i) *If $w(l) \notin \{j, j + 1\}$, then μ is a bottleneck matching for G' .*
- (ii) *If $w(l) \in \{j, j + 1\}$ and $G'(j)$ does not have a complete matching, then μ is a bottleneck matching for G' .*
- (iii) *If $w(l) \in \{j, j + 1\}$ and $G'(j)$ has a complete matching ν , then ν is a bottleneck matching for G' .*

Proof. Note first that $G(i) = G'(i)$ for all $i \neq j$ and recall the characterization of bottleneck assignments in Proposition 4.4-(ii).

(i): Let $w(l) = s \notin \{j, j + 1\}$. Since μ is a bottleneck matching for G , the graph $G(s)$ has a complete matching and $G(r)$ does not, for any $r < s$. Hence, the graph $G'(s) = G(s)$ has a complete matching and $G'(s - 1) = G(s - 1)$ does not. In addition, $\mu \subseteq G'(s)$, which ensures that μ is indeed a bottleneck matching for G' .

(ii): Since we assumed that $G'(j)$ has no complete matching, any complete matching in $G'(j + 1)$ is a bottleneck matching for G' . The matching μ is contained in $G'(j + 1) = G(j + 1)$ because we assumed $w(l) \in \{j, j + 1\}$.

(iii): It is clear that $G'(i) = G(i)$ for all $i < j$ and, hence, it does not have a complete matching. Since $\nu \subseteq G'(j)$, it is a bottleneck matching for G' . \square

As we have seen in Section 2, several edges in $A \times B$ can have the same length wherever the point set B is translated. However, this happens if and only if all such edges are between points $b \in B$ and $a \in A$ with the same vector $b - a$. In order to also handle point sets in this special position, we introduce the following equivalence relation.

Definition 4.8. Let $A, B \subset \mathbb{R}^2$ be finite point sets. Two edges $ab, a'b' \in A \times B$ are said to be *equivalent* if $b - a = b' - a'$.

Note that any two equivalent edges match distinct elements of B and, hence, the size of every equivalence class is at most k .


```

Input: planar points sets  $A, B \subset \mathbb{R}^2$  with  $k = |B| \leq |A| = n$ 
Output: labeled bottleneck diagram for  $A$  and  $B$ 
1 sort  $\{b - a : b \in B, a \in A\}$  lexicographically; group in equivalence
  classes
2 foreach  $b \in B$  do
3   compute  $(\leq k)$ -level of arrangement for  $S(b) = A - b$ 
4   foreach used bisector  $h = h(a, b, a', b)$  do
5      $\text{EdgePairs}(h) \leftarrow \text{EdgePairs}(h) \cup \{(ab, a'b)\}$ 
6   foreach  $b' \in B \setminus \{b\}$  do
7     compute  $(\leq 2k - 1)$ -level for  $S(b, b') = (A - b) \cup (A - b')$ 
8     foreach used bisector  $h = h(a, b, a', b')$  do
9        $\text{EdgePairs}(h) \leftarrow \text{EdgePairs}(h) \cup \{(ab, a'b')\}$ 
10  $\mathcal{L}(A, B) \leftarrow$  arrangement of the used bisectors
11  $C \leftarrow$  a cell of  $\mathcal{L}(A, B)$ ;  $t \leftarrow$  an interior point of  $C$ 
12  $S \leftarrow$  sorted list of  $\{\|b + t - a\| : b \in B, a \in A\}$ 
13  $G \leftarrow$  graph with edges  $Z(C)$  weighted by their order  $w$  in  $S$ 
14 label  $C$  with  $\mu \leftarrow$  bottleneck matching in  $G$  with respect to  $w$ 
15 while there is an unprocessed neighboring cell  $D$  of  $C$  do
16    $h \leftarrow D \cap C$ 
17   foreach  $(ab, a'b) \in \text{EdgePairs}(h)$  do
18     if  $(ab \in G, a'b \notin G)$  or  $(ab \notin G, a'b \in G)$  then
19        $G \leftarrow G \setminus \{\text{contained edge}\} \cup \{\text{non-contained edge}\}$ 
20        $l \leftarrow$  longest edge of  $\mu$ 
21        $\mu \leftarrow$  augmented matching  $\mu \setminus \{\text{contained edge}\}$  in  $G(w(l))$ 
22     if  $ab \in G$  and  $a'b \in G$  then
23       swap  $w(ab)$  and  $w(a'b)$  in  $G$ 
24   foreach  $(ab, a'b') \in \text{EdgePairs}(h)$  with  $b \neq b'$  do
25     swap  $w(ab)$  and  $w(a'b')$  in  $G$ 
26      $l \leftarrow$  longest edge of  $\mu$ ;  $j \leftarrow \min\{w(ab), w(a'b')\}$ 
27     if  $w(l) \in \{j, j + 1\}$  and  $G(j)$  has a complete matching  $\nu$  then
28        $\mu \leftarrow \nu$ 
29   label  $D$  with  $\mu$ 
30    $C \leftarrow D$ 

```

Algorithm 1: LabeledBottleneckDiagram(A, B)

Theorem 4.9. *Let $A, B \subset \mathbb{R}^2$ be with $k = |B| \leq |A| = n$. A labeled bottleneck diagram of A and B can be computed in $O(n^2 k^8)$ time, and a labeled lex-bottleneck diagram in $O(n^2 k^{10})$ time.*

Proof. We construct the diagrams by labeling the cells (and faces) of $\mathcal{L}(A, B)$ with a (lex-)bottleneck matching. A naive algorithm to do this would compute such a matching from scratch in every cell. However, we can maintain a bottleneck matching during a traversal of $\mathcal{L}(A, B)$, improving the time complexity of the algorithm. Unfortunately, this is not the case for the lex-bottleneck diagram, for which the best algorithm we know recomputes

(most of) the matching in a number of faces of $\mathcal{L}(A, B)$ that we are not able to bound away from its total complexity.

We detail first the algorithm to construct a bottleneck labeling. A pseudocode description is provided in Algorithm 1. We start by grouping the edges $b - a$, for $b \in B$ and $a \in A$, into equivalence classes, i.e., groups of edges that have the same length for any fixed translation, as defined in Definition 4.8. The involved sorting can be done in $O(nk \log n)$ time.

Then, we construct the line arrangement $\mathcal{L}(A, B)$ as described in Lemma 4.6, and we remember the involved edges for every used bisector. Note however that a bisecting line might be selected several times during this process. That is why we record a list of all the pairs of edges $(ab, a'b')$ that induce such a bisector (see Lines 5 and 9). As a consequence of Proposition 3.8, every such pair of edges inducing a fixed bisector of $S(b)$ is counted by Lemma 3.9. That is, if $h(a_1, b, a_2, b) = h(a_3, b, a_4, b)$, the first pair is counted as an order- j_1 edge and the second pair as an order- j_2 edge of $S(b)$ with $j_1 \neq j_2$. More precisely, since we could infinitesimally perturb the points in $S(b)$ such that $h(a_1, b, a_2, b) \neq h(a_3, b, a_4, b)$ for any choice of different points $a_1, a_2, a_3, a_4 \in A$ without altering the level of the edges inducing them, the bound in Lemma 3.9 counts already all the pairs inducing the same bisector. Every point in $S(b, b')$, for $b, b' \in B$ with $b \neq b'$, can correspond to two equivalent edges. Hence, for each “double” point inducing a bisector h , we add the corresponding additional pair of edges to the list $\text{EdgePairs}(h)$ (see Line 9). Therefore, the total number of edge pairs associated to bisectors is only a constant factor bigger than the bound on the total number of bisectors defining $\mathcal{L}(A, B)$. This shows that there is no overhead in the running time to record the edge pairs and so the Lines 2 to 10 take $O(n^2 k^6)$ time.

The final step of the algorithm is to traverse the just constructed arrangement $\mathcal{L}(A, B)$ and to find a bottleneck matching for each of its cells. We first initialize the traversal in an arbitrarily chosen cell C of $\mathcal{L}(A, B)$. To this end, we pick a point t interior to C , e.g., the centroid of its vertices, and we sort the values $\|b + t - a\|$ choosing one representative edge $ab \in A \times B$ from every equivalence class. We initialize also a graph G with the k^2 edges of $Z(C) = \bigcup_{b \in B} E_b(t)$, since we know by the graph theoretic version of Lemma 3.11-(ii) that it contains a, and in fact every, bottleneck label for C . Moreover, we construct the weight function $w : Z(C) \rightarrow [k^2]$ for G representing the order of the lengths of the edges of $Z(C)$ in the relative interior of C . We then find in Line 14 a bottleneck matching μ in G in $O(k^2 \sqrt{k \log k})$ time using the Gabow-Tarjan algorithm introduced in Section 4.1.

After this initialization, we now traverse $\mathcal{L}(A, B)$ while maintaining the graph G such that for the current cell D it has edges $Z(D)$ and weight function $w : Z(D) \rightarrow [k^2]$ encoding the relative lengths of these edges. This is done in Lines 14 to 30. Let h be the current bisector to be crossed. We first consider the edge pairs of the type $(ab, a'b)$ for h . If none of the edges belong to G , there is nothing to do. If both edges belong to G , Proposition 3.8 ensures that they have consecutive weights which need to be swapped during

the crossing of the bisector. The current bottleneck matching μ is not affected. If exactly one of the edges, say $a'b$, does not belong to the graph G , we include it and remove ab , because it is not part of the candidate set anymore. By Proposition 3.8, the weight of the new edge is the same as the old one. In addition, if the removed edge ab belongs to the current bottleneck matching μ with longest edge l , then a new matching having longest edge of weight $w(l)$ can be found in the updated graph. In other words, a pair involving edges incident to the same $b \in B$ may change the candidate set replacing an edge by another edge with the same weight, but the weight of the longest edge of the bottleneck matchings remains invariant. In order to find the new bottleneck matching, it is enough to augment $\mu \setminus \{ab\}$ in the graph $G(w(l))$. Such an augmenting path is guaranteed to exist by Lemma 4.2 because $G(w(l))$ has a complete matching.

Consider now all the edge pairs of the type $(ab, a'b')$, with $b \neq b'$, in the list of h . If at least one of the edges is not contained in G , there is nothing to do. Otherwise, the weights of the edges need to be swapped as in the previous case. The bottleneck matching is updated according to the rules described in Lemma 4.7, where we might need to test if $G'(j)$, in the notation of Lemma 4.7, has a complete matching. If m edges had weight j in G , the matching $\tau = \mu \cap G'(j)$ can have up to m exposed vertices. We search in $O(k^2)$ time for an augmenting path for τ in $G'(j)$. If there is one, we augment the matching and search again. If there is no augmenting path, the matching is complete by Lemma 4.2. Therefore, we can decide whether $G'(j)$ has a complete matching and, if not, find one performing at most as many augmentations as there are pairs in the list of h . After handling all the edge pairs of h , the resulting graph contains the candidate set of edges $Z(D)$ for the new cell D and w expresses the relative order of their lengths. Thus, the bottleneck matching we obtained for the last weighted graph is guaranteed to be a bottleneck matching for any point interior to D .

The number of graph and matching updates performed during the traversal is bounded by the number of edges that $\mathcal{L}(A, B)$ would have if we replace each bisector associated to s edge pairs by s infinitesimally-separated lines parallel to it. As argued before, the number of lines supporting edges of this arrangement would still be $O(nk^3)$ and, thus, the complexity of this virtual arrangement is $O(n^2k^6)$. Together with the required time $O(k^2)$ for updating the bottleneck matching in each cell, this implies the claimed running time $O(n^2k^8)$ for the whole algorithm.

In order to construct a lex-bottleneck labeling, we maintain the weighted graph as in the bottleneck case. We apply the algorithm described in [21] in every face of $\mathcal{L}(A, B)$, after updating the weight function w to indicate the ties that are active in the current face. This computes a lexicographic bottleneck matching in $O(k^4)$ time. \square

In higher dimensions, we do not have a good bound on the complexity of the reduced arrangement $\mathcal{L}(A, B)$. Nevertheless, the previous proof can be adapted using $\mathcal{H}(A, B)$ instead, leading via Corollary 3.6 to the following result.

Theorem 4.10. *Let $A, B \subset \mathbb{R}^d$ be point sets with $k = |B| \leq |A| = n$. A labeled bottleneck diagram of A and B can be computed in $O(n^{2d}k^{2d+2})$ time, and a labeled lex-bottleneck diagram can be computed in $O(n^{2d}k^{2d+4})$ time.*

5. APPLICATIONS

In this section, we explore some of the applications of the bottleneck diagrams studied before. An obvious application is to solve the bottleneck partial-matching problem under translation. Furthermore, bottleneck partial-matching Voronoi diagrams serve as a data structure for dynamically querying for locally optimal bottleneck matchings. Two such situations are described below.

5.1. Solving the Bottleneck Partial-Matching Problem. Let $A, B \subset \mathbb{R}^2$ be two point sets with $k = |B| \leq |A| = n$. We are interested here in finding a matching $\sigma : B \hookrightarrow A$ such that

$$f_\sigma(t^*) = \min_{t \in \mathbb{R}^2} \mathcal{E}(t), \text{ for some } t^* \in \mathbb{R}^2,$$

using the notation from Definition 2.1. Such a matching is called an *optimal bottleneck matching under translations*. The basic idea to find such a matching is to traverse a labeled bottleneck diagram and compute the optimal value of the cost function in every convex cell of the diagram separately. Clearly, this procedure also admits to report a corresponding translation t^* for which the minimum is attained.

Theorem 5.1. *Let $A, B \subset \mathbb{R}^2$ be sets of k and n points. An optimal bottleneck matching for A and B under translations can be found in $O(n^2k^8)$ time.*

Proof. We construct $\mathcal{L}(A, B)$ and a bottleneck labeling for it in time $O(n^2k^8)$ as described in the proof of Theorem 4.9. We traverse this arrangement and optimize in every (convex) cell C the value $f(t) = \|b + t - a\|^2$ over all translations $t \in C$, maintaining the minimum throughout the diagram. Here, ab is the longest edge of the bottleneck matching given by the label of the current cell C . More precisely, let $t_0 = a - b$ be the translation attaining the global minimum of the function $f(t)$. If $t_0 \in C$, obviously $f(t_0) = 0$ is the minimum of f in C and in fact, a global minimum as well. Otherwise, the minimum is attained in the point of C closest to t_0 . Such a point must be either a vertex of C or the orthogonal projection of t_0 onto an edge of C . In addition, if t_0 is the minimum of $f(t)$ and t_1 is the minimum of the corresponding function for a neighboring cell D sharing the edge e with C , then the projection of t_0 onto e coincides with the projection of t_1 onto e . We can thus calculate the minimum examining once every vertex of the diagram and at most one candidate point for every edge. Thus, the total time needed to perform the mentioned optimization in every cell is proportional to the complexity of the diagram. \square

5.2. Computing a Bottleneck Path. We consider now the problem of finding a motion for B from an initial position to a final position such that the maximum bottleneck value (as defined in Definition 2.1) attained during the motion is minimized.

Definition 5.2. The *bottleneck value* of a curve $\gamma : [0, 1] \rightarrow \mathbb{R}^2$ with respect to point sets $A, B \subset \mathbb{R}^2$ with $k = |B| \leq |A| = n$ is

$$F(\gamma) = \max_{s \in [0, 1]} \mathcal{E}(\gamma(s)) = \max_{s \in [0, 1]} \min_{\sigma: B \hookrightarrow A} \max_{b \in B} \|b + \gamma(s) - \sigma(b)\|^2.$$

The curve γ is called a *bottleneck path* if $F(\gamma) \leq F(\varphi)$ for every other curve $\varphi : [0, 1] \rightarrow \mathbb{R}^2$ with $\varphi(0) = \gamma(0)$ and $\varphi(1) = \gamma(1)$.

A bottleneck path between two positions can be useful in motion planning where the points of A represent fixed anchor points and the points of B represent the position of articulations of a moving robot. The dual graph of the arrangement $\mathcal{L}(A, B)$ contains the necessary information to compute a bottleneck path from any initial position to any final position.

Definition 5.3. The *bottleneck graph* of two finite point sets $A, B \subset \mathbb{R}^2$ is the weighted graph $\mathcal{L}(A, B)^*$ dual to $\mathcal{L}(A, B)$, where an edge e^* of the graph dual to an edge e of $\mathcal{L}(A, B)$ has weight $\min_{t \in e} \mathcal{E}(t)$.

Via the bottleneck graph we can now characterize the existence of a path of given bottleneck value.

Lemma 5.4. Let $t_0, t_1 \in \mathbb{R}^2$ and $\delta \in \mathbb{R}$. Let C_0 and C_1 be cells of $\mathcal{L}(A, B)$ such that $t_0 \in C_0$ and $t_1 \in C_1$. There is a path with bottleneck value at most δ from t_0 to t_1 if and only if $\mathcal{E}(t_0), \mathcal{E}(t_1) \leq \delta$ and there is a path from C_0^* to C_1^* in $\mathcal{L}(A, B)^*$ whose longest edge has weight at most δ .

Proof. Observe that in every cell of $\mathcal{L}(A, B)$ there is a bottleneck matching whose cost coincides with \mathcal{E} in the cell. By definition, \mathcal{E} is a convex function in every such (convex) cell. Hence, assuming that $C_0 = C_1$, the line segment joining t_0 and t_1 has bottleneck value $\max\{\mathcal{E}(t_0), \mathcal{E}(t_1)\}$ and no path can attain a smaller value. We assume now that $C_0 \neq C_1$ and let γ be any path from t_0 to t_1 . We can replace each of the connected arcs of γ entering a cell C of $\mathcal{L}(A, B)$ in a point t_{in} and leaving it in a point t_{out} by the line segment joining these two points without increasing the bottleneck value of the path. Again, we do not increase the bottleneck value of the path when we substitute this line segment by the one joining the points t_{in}^* and t_{out}^* , where t_{in}^* is the point with minimum bottleneck value on the edge of C that contains t_{in} , and t_{out}^* is the one attaining the minimum value on the edge containing t_{out} . Similarly, the parts of the path in C_0 and C_1 , starting at t_0 and ending at t_1 , respectively, can be replaced with the line segment from t_0 (or t_1) to the point attaining the minimum of \mathcal{E} on whichever edge of C_0 (or C_1) the path crosses first (or last).

The previous observations imply that a bottleneck path is among the polygonal paths whose vertices (except for t_0 and t_1) lie on the minima of \mathcal{E} along edges of $\mathcal{L}(A, B)$. The bottleneck value of such a path is the maximum of the weights of the edges in the corresponding path in $\mathcal{L}(A, B)^*$ and the values $\mathcal{E}(t_0)$ and $\mathcal{E}(t_1)$. \square

Theorem 5.5. Given $t_0, t_1 \in \mathbb{R}^2$, a bottleneck path from t_0 to t_1 with respect to $A, B \subset \mathbb{R}^2$ with $k = |B| \leq |A| = n$ can be computed in time $O(n^2 k^6 (k^2 + \log n))$.

Proof. We first compute the arrangement $\mathcal{L}(A, B)$ and the associated bottleneck graph $\mathcal{L}(A, B)^*$ in time $O(n^2k^8)$ by Theorem 4.9. The number of edges and vertices of $\mathcal{L}(A, B)^*$ is $O(n^2k^6)$ due to Theorem 3.12 and the weights of its edges are all nonnegative. Therefore, the path with minimum bottleneck value in the graph can be found in $O(n^2k^6 \log n)$ time via the implementation of Dijkstra's algorithm using heaps. By Lemma 5.4, the associated polygonal path is guaranteed to be a bottleneck path from t_0 to t_1 . \square

5.3. Finding the Cover Radius of a Convex Polygon. As a third application of our bottleneck diagrams, we investigate a covering problem. Given a pair of finite point sets $A, B \subset \mathbb{R}^2$, and a convex polygon $Q \subset \mathbb{R}^2$, we want to determine the minimal $\delta \in \mathbb{R}$ such that for any position of the point set B in Q there is a matching whose bottleneck value is at most δ . We can think of the points in A as antennas equipped with disks of radius δ modeling the region on which they provide signal. The point set B can be thought of as a robot that moves in Q and needs to connect each of its points to a different antenna (for instance, to learn its position). The target is to minimize the power consumed by the antennas while ensuring that the robot can move in Q having a one-to-one connection for its receivers.

Definition 5.6. The *cover radius* of a convex polygon Q with respect to finite point sets $A, B \subset \mathbb{R}^2$ is the maximum bottleneck value among all $t \in \mathbb{R}^2$ such that $B + t \subset Q$.

Theorem 5.7. Let Q be a convex polygon with m vertices. The cover radius of Q with respect to $A, B \subset \mathbb{R}^2$ with $k = |B| \leq |A| = n$ can be computed in time $O(n^2k^8 + (n^2k^6 + m) \log(n + m))$.

Proof. We start by computing $\mathcal{L}(A, B)$ and a bottleneck labeling as indicated in Theorem 4.9. Note that the set of translations $t \in \mathbb{R}^2$ for which $B + t \subset Q$ is a convex polygon \hat{Q} , which is the intersection of k translated copies of Q . This polygon is indeed given by the m linear inequalities obtained by imposing that the extreme point of B in the direction orthogonal to an edge of Q is on the right hand side of the corresponding edge. This polytope can be computed easily in $O(k \log k + m \log m)$ time and has at most m edges. Then, we compute the overlay of the boundary of \hat{Q} and $\mathcal{L}(A, B)$. Note that every edge of $\mathcal{L}(A, B)$ can intersect the boundary of \hat{Q} in at most two points. Thus, the number of vertices of the overlay is $O(n^2k^6 + m)$ and, hence, it can be computed in $O((n^2k^6 + m) \log(n + m))$ time using the techniques described in [10]. The next step involves traversing the overlay and maintaining the maximum of \mathcal{E} in every cell. Since the function is convex, the maximum can be calculated as the maximum of the values attained at the vertices of the overlay. \square

6. CONCLUSION AND OPEN PROBLEMS

In this work, we introduced and investigated Voronoi-type diagrams suited for the study of the bottleneck partial-matching problem under translations. As our main results, we obtained low complexity bounds on these diagrams and devised efficient algorithms for their construction that allowed to solve

the matching problem both in arbitrary dimension and in its lexicographic variant.

We have seen that the complexity bound of $O(n^2k^6)$ in Theorem 3.12 is sharp in terms of the parameter n , i.e., the size of the bigger point set, but we do not know whether it might be improved with respect to k . Since any improvement of this kind that comes with a speed-up in the construction of the (lex-)bottleneck diagram translates into a better bound on the running time of our algorithms for the partial-matching problems, we consider this as an interesting open problem for future research. In fact, we believe that the lower bound in Proposition 3.13 is optimal.

Conjecture 6.1. *For any pair of point sets $A, B \subset \mathbb{R}^2$ with $k = |B| \leq |A| = n$, there is a lex-bottleneck diagram of complexity $O(n^2k^2)$.*

REFERENCES

- [1] H. Alt, N. Blum, K. Mehlhorn, and M. Paul. Computing a maximum cardinality matching in a bipartite graph in time $O(n^{1.5}\sqrt{m/\log n})$. *Information Processing Letters*, 37(4):237–240, Feb. 1991.
- [2] H. Alt, K. Mehlhorn, H. Wagener, and E. Welzl. Congruence, similarity, and symmetries of geometric objects. *Discrete & Computational Geometry*, 3(1):237–256, 1988.
- [3] R. Ben-Avraham, M. Henze, O. E. Raz, R. Jaume, M. Sharir, B. Keszegh, and I. Tubis. Minimum Partial-Matching and Hausdorff RMS-Distance under Translation: Combinatorics and Algorithms. *submitted full version*. <http://arxiv.org/abs/1411.7273>.
- [4] R. Ben-Avraham, M. Henze, O. E. Raz, R. Jaume, M. Sharir, B. Keszegh, and I. Tubis. Minimum Partial-Matching and Hausdorff RMS-Distance under Translation: Combinatorics and Algorithms. In A. S. Schulz and D. Wagner, editors, *Proc. 22nd Annual European Symposium on Algorithms*, volume 8737 of *LNCS*, pages 100–111, 2014.
- [5] C. Berge. Two theorems in graph theory. In *Proc. of the National Academy of Sciences of the United States of America*, volume 43, pages 842–844, 1957.
- [6] R. E. Burkard, M. Dell’Amico, and S. Martello. *Assignment Problems*. Society for Industrial and Applied Mathematics, 2009.
- [7] R. E. Burkard and F. Rendl. Lexicographic bottleneck problems. *Operations Research Letters*, 10(5):303–308, 1991.
- [8] T. M. Chan. Random sampling, halfspace range reporting, and construction of $(\leq k)$ -levels in three dimensions. *SIAM Journal on Computing*, pages 586–595, 1999.
- [9] B. Chazelle, L. Guibas, and D. Lee. The power of geometric duality. *BIT Numerical Mathematics*, 25(1):76–90, 1985.
- [10] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag TELOS, Santa Clara, CA, USA, 3rd ed. edition, 2008.
- [11] F. Della Croce, V. T. Paschos, and A. Tsoukias. An improved general procedure for lexicographic bottleneck problems. *Operations Research Letters*, 24(4):187–194, 1999.
- [12] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, Berlin, Heidelberg, 1987.
- [13] H. Edelsbrunner and L. J. Guibas. Topologically sweeping an arrangement. *Journal of Computer and System Sciences*, 38(1):165–194, 1989.
- [14] A. Efrat, A. Itai, and M. J. Katz. Geometry helps in bottleneck matching and related problems. *Algorithmica*, 31(1):1–28, 2001.
- [15] H. N. Gabow and R. E. Tarjan. Algorithms for two bottleneck optimization problems. *Journal of Algorithms*, 9(3):411–417, Sept. 1988.

- [16] M. Henze and R. Jaume. Bottleneck Partial-Matching Voronoi Diagrams and Applications. In *Proc. 25th International Symposium on Algorithms and Computation*, LNCS, 2014.
- [17] M. Henze, R. Jaume, and B. Keszegh. On the complexity of the partial least-squares matching Voronoi diagram. In *Proc. 29th European Workshop on Computational Geometry*, pages 193–196, 2013.
- [18] J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, 2(4):225–231, 1973.
- [19] A. P. Punnen and K. Nair. Improved complexity bound for the maximum cardinality bottleneck bipartite matching problem. *Discrete Applied Mathematics*, 55(1):91–93, 1994.
- [20] G. Rote. Partial least-squares point matching under translations. In *Proc. 26th European Workshop on Computational Geometry*, pages 249–251, 2010.
- [21] P. T. Sokkalingam and Y. P. Aneja. Lexicographic bottleneck combinatorial problems. *Operations Research Letters*, 23(1–2):27–33, 1998.

APPENDIX A. k -SENSITIVE ANALYSIS OF PREVIOUS ALGORITHMS

The following k -sensitive analysis of previous algorithms for the bottleneck partial-matching problem under translation is not intended to be a complete account of all involved details. Rather we focus on the crucial ideas and adjustments needed in order to make these algorithms work in the general situation.

A.1. The Algorithm of Alt, Mehlhorn, Wagener & Welzl [2]. The first algorithm was introduced in [2], where the authors prove a running time of $O(n^6 \log n)$ for the balanced situation. Their approach consists of two steps and can be adjusted toward the unbalanced case as follows. First, for every choice of $a_1, a_2, a_3 \in A$ and $b_1, b_2, b_3 \in B$ they define $\varepsilon(b_1, a_1, b_2, a_2, b_3, a_3)$ to be the minimum $\varepsilon \in \mathbb{R}$ such that there is a translation placing b_i into an ε -neighborhood of a_i , for all $i = 1, 2, 3$. They claim that the bottleneck distance under translations is attained by one of these $O(n^6)$ values, and they compute every such value in constant time. In the unbalanced case, the number of values and the time for its computation is then $O(n^3 k^3)$. They sort these values into an array \mathfrak{E} and perform a binary search, testing for every $\varepsilon \in \mathfrak{E}$ whether there is a bottleneck matching under translations having cost ε . In order to do that, they assume that $\|b + t - a\|^2 = \varepsilon^2$, for a fixed pair $ab \in A \times B$, which restricts the set of candidate translations to a circle. They parametrize the circle by polar coordinates and compute the set of angles $\alpha(a', b')$ for which b' lies in an ε -neighborhood of a' , for all $a'b' \in A \times B$. The computation of such intervals on the circle is not trivial and requires some careful observations that carry over to the unbalanced case without modification. The arrangement in this circle induced by the circular intervals α can then be computed by sorting their endpoints. In every interval, the authors construct the bipartite graph whose edges are shorter than ε and they look for a maximum matching in it. This is done by computing the graph for an arbitrary initial point and traversing the circular arrangement, adding or deleting at each interval the corresponding edge or edges. If edges are only added, nothing is left to be done. If some edges are deleted leaving some points of B unmatched, suitable augmenting-path computations need to be performed in order to decide whether there is a maximum matching in the next interval.

The construction of each of the arrangements in the $O(nk)$ circles for a fixed value $\varepsilon \in \mathfrak{E}$ can be done in $O(nk \log n)$. The traversal requires $O(nk)$ time to construct the initial graph, $O(nk + k^{5/2})$ to prune non-relevant edges and obtain a maximum matching for the initial graph using the Hopcroft-Karp algorithm, and $O(nk)$ time per cell to compute the augmenting path to update the maximum matching. The updates of the bipartite graph require constant time for each edge. Thus, the total time required for each $\varepsilon \in \mathfrak{E}$ is $O(nk \cdot (nk \log n + nk + k^{5/2} + nk \cdot nk)) = O(n^3 k^3)$. Since sorting the values in \mathfrak{E} is done in $O(n^3 k^3 \log n)$ time, the whole algorithm runs in time $O(n^3 k^3 \log n)$.

A.2. The Algorithm of Efrat, Itai & Katz [14]. An improved algorithm with a running time of $O(n^5 \log^2 n)$ in the balanced case has been described in [14]. This improvement requires the use of a non-trivial data structure and parametric search techniques combined with sorting networks. The data structure is used to create an oracle that, given $\varepsilon \in \mathfrak{E}$ in the above notation, answers whether the bottleneck distance of A and B under translations is at most ε . The corresponding oracle in [2] runs in $O(n^6)$ time and, according to the preceding analysis, can be adapted to run in $O(n^3 k^3)$ in the unbalanced case. A gain in the computation of augmenting paths enables the authors in [14] to improve this oracle to run in $O(n^5 \log n)$ time. This is done by utilizing a data structure based on constructing a layered graph from the matching after each possible edge deletion. Adapting the analysis for the unbalanced case leads to an augmenting paths computation in time $O(k \log n)$ instead of the $O(nk)$ by standard techniques. Therefore, the oracle runs in $O(n^2 k^3 \log n)$ time. However, the time to sort the $O(n^3 k^3)$ values into \mathfrak{E} dominates the running time and thus prevents the new algorithm to improve upon the simpler one in [2]. This problem is solved by using, instead of the binary search, an adaptation of the parametric search technique due to Cole in order to avoid the construction of \mathfrak{E} , and hence reducing the number of oracle calls to $O(\log n)$. Therefore, the final running time for the general unbalanced situation is $O(n^2 k^3 \log^2 n)$.

INSTITUT FÜR INFORMATIK, FREIE UNIVERSITÄT BERLIN, TAKUSTRASSE 9, 14195
BERLIN, GERMANY

E-mail address: matthias.henze@fu-berlin.de

INSTITUT FÜR INFORMATIK, FREIE UNIVERSITÄT BERLIN, TAKUSTRASSE 9, 14195
BERLIN, GERMANY

E-mail address: jaume@inf.fu-berlin.de